

CONSISTENCY OF SPARSE PCA IN HIGH DIMENSION, LOW SAMPLE SIZE CONTEXTS

BY DAN SHEN^{*,†}, HAIPENG SHEN[‡] AND J. S. MARRON[§]

University of North Carolina at Chapel Hill

Sparse Principal Component Analysis (PCA) methods are efficient tools to reduce the dimension (or the number of variables) of complex data. Sparse principal components (PCs) are easier to interpret than conventional PCs, because most loadings are zero. We study the asymptotic properties of these sparse PC directions for scenarios with fixed sample size and increasing dimension (i.e. High Dimension, Low Sample Size (HDLSS)). Under the previously studied spike covariance assumption, we show that Sparse PCA remains consistent under the same large spike condition that was previously established for conventional PCA. Under a broad range of small spike conditions, we find a large set of sparsity assumptions where Sparse PCA is consistent, but PCA is strongly inconsistent. The boundaries of the consistent region are clarified using an oracle result.

1. Introduction. Principal Component Analysis (PCA) is an important visualization and dimension reduction tool for High Dimension, Low Sample Size (HDLSS) data. However, the linear combinations found by PCA typically will involve all the variables, with non-zero loadings, which can be challenging to interpret. In order to overcome this weakness of PCA, we will study sparse PCA methods that generate sparse principal components (PCs), i.e. PCs with only a few non-zero loadings. Several sparse PCA methods have been proposed to facilitate the interpretation of HDLSS data, see for example Zou, Hastie and Tibshirani (2006) [?], Shen and Huang (2008) [?], Leng and Wang (2009) [?], Witten, Tibshirani and Hastie (2009) [?], Johnstone and Lu (2009) [?], Amini and Wainwright (2009) [?], and Ma (2010) [?].

This paper studies the HDLSS asymptotic properties of sparse PCA. HDLSS asymptotics are based on the limit, as the dimension $d \rightarrow \infty$, with the sample size n fixed, as originally studied by Hall, Marron and

^{*}Corresponding Author

[†]Partially supported by NSF grant DMS-0854908

[‡]Partially supported by NSF grants DMS-0606577 and CMMI-0800575

[§]Partially supported by NSF grants DMS-0606577 and DMS-0854908

AMS 2000 subject classifications: Primary 62H25; secondary 62F12

Keywords and phrases: Sparse PCA, High Dimension, Low Sample Size, Consistency

Neeman (2005) [?] and Ahn et al. (2007) [?]. Theoretical properties of sparse PCA have been studied before under different asymptotic frameworks. Leng and Wang (2009) [?] used the *adaptive lasso* penalty of Zou, Hastie and Tibshirani (2006) [?] to introduce sparse loadings, and established some consistency result for selecting non-zero loadings when the sample size $n \rightarrow \infty$, with the dimension d fixed. Johnstone and Lu (2009) [?] considered a single-component spiked covariance model (originally proposed by Johnstone (2001) [?]) and showed that conventional PCA is consistent if and only if $d(n)/n \rightarrow 0$; furthermore, under the condition $\log(d \vee n)/n \rightarrow 0$, they proved consistency of PCA performed on a subset of variables with largest sample variance. Amini and Wainwright (2009) [?] considered the same single-component spiked model, and further restricted the maximal eigenvector to have k non-zero entries; they studied the thresholding subset PCA procedure of Johnstone and Lu [?] and the sparsePCA procedure of d’Aspremont et al. (2007) [?], and explored conditions on the triplet (n, d, k) under which each procedure can recover the support set of the sparse eigenvector with probability one. Paul and Johnstone [?] developed the augmented sparse PCA procedure along with its optimal rate of convergence property. Ma [?] proposed an iterative thresholding procedure for estimating principal subspaces that has nice theoretical properties.

Sparse PCA is primarily motivated by modern data sets of very high dimension; hence we prefer the statistical viewpoint of the High Dimension Low Sample Size (HDLSS) asymptotics. Note that this case of $d \rightarrow \infty$ with n fixed was not considered by Johnstone and Lu [?]. Conventional PCA was first studied using HDLSS asymptotics by Ahn et al. [?] and the most comprehensive current result is Jung and Marron (2009) [?]. The latter found conditions when the first several empirical PC directions would be *consistent* or *subspace consistent* with the corresponding population PC directions. This happens when the first several eigenvalues are large enough, compared with the rest of the eigenvalues of the population covariance matrix. Moreover, if the first few eigenvalues are not sufficiently large, all empirical PC directions will be *strongly inconsistent* with their population counterparts in the sense that the angle between them will converge to 90 degrees.

The main contribution of this paper is an exploration of conditions where conventional PCA is strongly inconsistent (for scenarios with relatively small population eigenvalues), yet sparse PCA methods are consistent. Furthermore, the mathematical boundaries of the sparse PCA consistency are established by showing strong inconsistency, for even an oracle version of sparse PCA, beyond the consistent region. Similar to Johnstone and Lu (2009) [?] and Amini and Wainwright (2009) [?], we focus on the single component

spiked covariance model. Our results depend on a *spike index*, α , defined below in the context of Example 1.1, which measures the dominance of the first eigenvalue, and on a *sparsity index*, β , defined also in Example 1.1, which measures the number of non-zero entries of the first population eigenvector. For illustration purposes, we simplify the consistency and strong inconsistency results for the exemplary model considered in Example 1.1, and summarize them below as functions of α and β in Figure 1:

- **Previous Results (dark grey rectangle):** Jung and Marron (2009) [?] showed that the first empirical eigenvector is consistent with the first population eigenvector when the spike index α is greater than 1.
- **Consistency (white triangle):** We will show that sparse PCA is consistent even when the spike index α is less than 1, as long as α is greater than the sparsity index β . This is done in Section 2 for a simple thresholding method and in Section 3 for the RSPCA method proposed by Shen and Huang (2008) [?].
- **Strong Inconsistency (black triangle):** In Section 4 we show that even an oracle sparse PCA procedure is strongly inconsistent with the first population eigenvector, when the spike index α is smaller than the sparsity index β .
- **Irrelevant Area (light grey rectangle):** The sparsity index β can not be larger than 1, hence the light grey rectangular area is irrelevant.

1.1. *Notation and Assumptions.* All quantities are indexed by the dimension d in this paper. However, when it will not lead to confusion, the subscript d will be omitted for convenience. Let the population covariance matrix be Σ_d . The eigen-decomposition of Σ_d is

$$\Sigma_d = U_d \Lambda_d U_d^T,$$

where Λ_d is the diagonal matrix of the population eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ and U_d is the matrix of corresponding population eigenvectors so that $U_d = [u_1, \dots, u_d]$.

Assume that X_1, \dots, X_n are random samples from a d -dimensional normal distribution $N(0, \Sigma_d)$. Denote the data matrix by $X_{(d)} = [X_1, \dots, X_n]_{d \times n}$ and the sample covariance matrix by $\hat{\Sigma}_d = n^{-1} X_{(d)} X_{(d)}^T$. Then, the sample covariance matrix $\hat{\Sigma}_d$ can be similarly decomposed as

$$\hat{\Sigma}_d = \hat{U}_d \hat{\Lambda}_d \hat{U}_d^T,$$

where $\hat{\Lambda}_d$ is the diagonal matrix of the sample eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ and \hat{U}_d is the matrix of the corresponding sample eigenvectors so that $\hat{U}_d = [\hat{u}_1, \dots, \hat{u}_d]$.




FIG 1. Consistent areas for PCA and sparse PCA, as a function of the spike index α and the sparsity index β , under the single component spiked model considered in Example 1.1. Conventional PCA is consistent only on the dark grey rectangle ($\alpha > 1$), while sparse PCA is also consistent on the white triangle ($0 \leq \beta < \alpha \leq 1$). In addition, an oracle sparse PCA procedure is strongly inconsistent on the black triangle ($0 \leq \beta < \alpha \leq 1$). The light grey rectangular area ($0 \leq \alpha < 1, \beta > 1$) is not considered because the sparsity index $\beta \leq 1$. The dots show the grid points studied in the simulation study of Section 4.

Let \bar{u}_i be any sample based estimator of u_i , e.g. $\bar{u}_i = \hat{u}_i$ for $i = 1, \dots, d$. Two important concepts from Jung and Marron (2009) [?] are:

- **Consistency:** The direction \bar{u}_i is *consistent* with its population counterpart u_i if

$$(1.1) \quad \text{Angle}(\bar{u}_i, u_i) \equiv \arccos(|\langle \bar{u}_i, u_i \rangle|) \xrightarrow{p} 0, \text{ as } d \rightarrow \infty,$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product between two vectors.

- **Strong Inconsistency:** The direction \bar{u}_i is *strongly inconsistent* with its population counterpart u_i if

$$\text{Angle}(\bar{u}_i, u_i) = \arccos(|\langle \bar{u}_i, u_i \rangle|) \xrightarrow{p} \frac{\pi}{2}, \text{ as } d \rightarrow \infty.$$

In addition, we consider another important concept in the current paper:

- **Consistency with convergence rate d^ι :** The direction \bar{u}_i is consistent with its population counterpart u_i with the convergence rate d^ι if $|\langle \bar{u}_i, u_i \rangle| = 1 + o_p(d^{-\iota})$, where the notation $G_d \equiv o_p(d^{-\iota})$ means that $d^\iota G_d \xrightarrow{p} 0$, as $d \rightarrow \infty$.

EXAMPLE 1.1. Assume that X_1, \dots, X_n are random sample vectors from a d -dimensional normal distribution $N(0, \Sigma_d)$, where the covariance matrix Σ_d has the eigenvalues as

$$\lambda_1 = d^\alpha, \lambda_2 = \dots = \lambda_d = 1, \alpha \geq 0.$$

This is a special case of the single component spike covariance Gaussian model considered before by, for example, Johnstone (2001) [?], Paul (2007) [?], Johnstone and Lu (2009) [?], Amini and Wainwright (2009) [?]. Without loss of generality (WLOG), we further assume that the first eigenvector u_1 is proportional to the following d -dimensional vector

$$\dot{u}_1 = (\overbrace{1, \dots, 1}^{\lfloor d^\beta \rfloor}, 0, \dots, 0)^T,$$

where $0 \leq \beta \leq 1$ and $\lfloor d^\beta \rfloor$ denotes the integer part of d^β . (In general the non-zero entries do not have to be the first $\lfloor d^\beta \rfloor$ elements, neither do they need to be equal.) If $\beta = 0$, the first population eigenvector becomes $u_1 = (1, 0, \dots, 0)^T$.

For the above model, Jung and Marron (2009) [?] showed that the first empirical eigenvector (the PC direction) \hat{u}_1 is consistent with u_1 when $\alpha > 1$; however for $\alpha < 1$, it is strongly inconsistent. Again, the main point of the current paper is an exploration of conditions under which sparse methods can lead to consistency when the spike index $\alpha \leq 1$, (recall that the first eigenvalue $\lambda_1 = d^\alpha$), by exploiting *sparsity*. Sparsity is quantified by the sparsity index β , where $\lfloor d^\beta \rfloor$ is the number of non-zero elements of the first eigenvector u_1 . Here we use the above simple example for intuitive illustration purposes, to highlight the key findings. More general single component spike models will be considered in Sections 2 to 4.

1.2. *Roadmap of the paper.* The organization of the rest of paper is as follows. For easy access to the main ideas, Section 2 first introduces a simple thresholding method to generate sparse PC directions. Section 2.1 shows the consistency of the sparse PC directions, obtained by this simple thresholding method. Section 3 then generalizes these ideas to a current sparse PCA method. In particular, we consider the sparse PCA method developed by

Shen and Huang (2008) [?], and build its connection to the simple thresholding method. We then establish the consistency of the sparse PCA method under the sparsity and small spike conditions where the conventional PCA is strongly inconsistent. Section 4 considers scenarios when the spike index α is smaller than the sparsity index β , and proves the strong inconsistency of an appropriate oracle PCA procedure. Section 5 reports some simulation results to illustrate both consistency and strong inconsistency of PCA and sparse PCA. Section 6 concludes the paper with some discussion of future work on extending consistency of sparse PCA to more general distributions. We point out that it is challenging to move beyond Gaussianity to get HDLSS consistency of sparse PCA. Section 7 contains the proofs of the theorems.

2. Consistency of a simple thresholding method for sparse PCA

in HDLSS. In Example 1.1, the first eigenvector of the sample covariance matrix \hat{u}_1 is strongly inconsistent with u_1 when $\alpha < 1$, because it attempts to estimate too many parameters. Sparse data analytic methods assume that many of these parameters are zero, which can allow greatly improved estimation of the first PC direction u_1 . Here, this issue is explored in the context of sparse PCA, where $u_1 = (1, 0, \dots, 0)^T$ is an extreme case. The sample covariance matrix based estimator, \hat{u}_1 , can be improved by exploiting the fact that u_1 has many zero elements.

A natural approach is a simple thresholding method where entries with small absolute values are replaced by zero. In HDLSS contexts, it is challenging to apply thresholding directly to the entries of \hat{u}_1 , because the number of them grows rapidly as $d \rightarrow \infty$, which naturally shrinks their magnitudes given that \hat{u}_1 is a unit vector. Thresholding is more conveniently formulated in terms of the *dual covariance matrix* as used by Jung, Sen and Marron (2010) [?].

Denote the dual sample covariance matrix by $S_d = \frac{1}{n} X_{(d)}^T X_{(d)}$ and the first dual eigenvector by \tilde{v}_1 . The sample eigenvector \hat{u}_1 is connected with the dual eigenvector \tilde{v}_1 through the following transformation,

$$(2.1) \quad \tilde{u}_1 = (\tilde{u}_{1,1}, \dots, \tilde{u}_{d,1})^T = X_{(d)} \tilde{v}_1,$$

and the sample estimate is then given by $\hat{u}_1 = \tilde{u}_1 / \|\tilde{u}_1\|$ [?].

Given a sequence of threshold values λ , define the thresholded entries as

$$(2.2) \quad \check{u}_{i,1} = \begin{cases} \tilde{u}_{i,1} & \text{if } |\tilde{u}_{i,1}| > \lambda, \\ 0 & \text{if } |\tilde{u}_{i,1}| \leq \lambda, \end{cases} \quad \text{for } i = 1, \dots, d.$$

Denote $\check{u}_1 = (\check{u}_{1,1}, \dots, \check{u}_{d,1})^T$ and normalize it to get the simple thresholding (ST) estimator $\hat{u}_1^{\text{ST}} = \check{u}_1 / \|\check{u}_1\|$.

For the model considered in Example 1.1, given an eigenvalue of strength $\alpha \in (0, 1)$, (recall $\lambda_1 = d^\alpha$ and \hat{u}_1 is strongly inconsistent), below we explore conditions on the threshold sequence λ under which the ST estimator \hat{u}_1^{ST} is in fact consistent with u_1 . First of all, the threshold λ can not be too large; otherwise all the entries will be zeroed out. It will be seen in Theorem 2.1 that a sufficient condition for this is $\lambda \leq d^{\frac{\gamma}{2}}$, where $\gamma \in (0, \alpha)$. Secondly, the threshold λ can not be too small, or pure noise terms will be included. A parallel sufficient condition is shown to be $\lambda \geq \log^\delta(d)$, where $\delta \in (\frac{1}{2}, \infty)$.

2.1. Consistency of the simple thresholding method. Below we formally establish conditions on the eigenvalues of the population covariance matrix Σ_d and the thresholding parameter λ , which give consistency of \hat{u}_1^{ST} to u_1 . All the technical proofs are provided in Section 7 and the supplement materials.

We begin with considering the extreme sparsity case $u_1 = (1, 0, \dots, 0)^T$. Suppose that $\lambda_1 \sim d^\alpha$, in the sense that $0 < c_1 \leq \underline{\lim}_{d \rightarrow \infty} \frac{\lambda_1}{d^\alpha} \leq \overline{\lim}_{d \rightarrow \infty} \frac{\lambda_1}{d^\alpha} \leq c_2$, where c_1 and c_2 are two constants. Similarly, assume $\sum_{i=2}^d \lambda_i \sim d$. As in Jung and Marron [?], denote the measure of sphericity as

$$\varepsilon \equiv \frac{\text{tr}^2(\Sigma_d)}{d \text{tr}(\Sigma_d^2)} = \frac{(\sum_{i=1}^d \lambda_i)^2}{d \sum_{i=1}^d \lambda_i^2},$$

and assume the ε -condition: $\varepsilon \gg \frac{1}{d}$, i.e

$$(2.3) \quad (d\varepsilon)^{-1} = \frac{\sum_{i=1}^d \lambda_i^2}{(\sum_{i=1}^d \lambda_i)^2} \rightarrow 0, \text{ as } d \rightarrow \infty.$$

Now we need to impose the following conditions on the eigenvalues:

- Assume that $\overline{\lim}_{d \rightarrow \infty} \frac{\lambda_1}{\sum_{i=2}^d \lambda_i} = c$, where c is a non-negative constant, and the ε -condition is satisfied. These conditions can guarantee that the dual matrix S_d has a limit. Hence the first dual eigenvector \hat{v}_1 will have a limit and it will then help build up the consistency of \hat{u}_1^{ST} .
- In addition, we need the second eigenvalue λ_2 to be an obvious distance away from the first eigenvalue λ_1 . If not, it will be hard to distinguish the first and second empirical eigenvectors as observed by Jung and Marron, among others. In that case the appropriate amount of thresholding on the first empirical eigenvector becomes unclear. Therefore, we assume that $\lambda_2 \sim d^\theta$, where $\theta < \alpha$.

THEOREM 2.1. *Suppose that X_1, \dots, X_n are random samples from a d -dimensional normal distribution $N(0, \Sigma_d)$ and the first population eigenvector $u_1 = (1, 0, \dots, 0)^T$. If the following conditions are satisfied:*

- (a) $\lambda_1 \sim d^\alpha$, $\lambda_2 \sim d^\theta$, and $\sum_{i=2}^d \lambda_i \sim d$, where $\theta \in [0, \alpha)$ and $\alpha \in (0, 1]$,
- (b) for a non-negative constant c , $\lim_{d \rightarrow \infty} \frac{\lambda_1}{\sum_{i=2}^d \lambda_i} = c$ and the ε -condition (2.3) is satisfied,
- (c) $\log^\delta(d) d^{\frac{\theta}{2}} \leq \lambda \leq d^{\frac{\gamma}{2}}$, where $\delta \in (\frac{1}{2}, \infty)$ and $\gamma \in (\theta, \alpha)$,

then the simple thresholding estimator \hat{u}_1^{ST} is consistent with u_1 .

In fact, $u_1 = (1, 0, \dots, 0)^T$ in Theorem 2.1 is a very extreme case. The following theorem considers the general case $u_1 = (u_{1,1}, \dots, u_{d,1})^T$, where only $\lfloor d^\beta \rfloor$ elements of u_1 are non-zero. WLOG, we assume that the first $\lfloor d^\beta \rfloor$ entries are non-zero just for notational convenience.

Define

$$(2.4) \quad Z_j \equiv (z_{1,j}, \dots, z_{d,j})^T = (X_j^T u_1, \dots, X_j^T u_d)^T, \quad j = 1, \dots, n.$$

We can show that Z_j are iid $N(0, \text{diag}\{\lambda_1, \dots, \lambda_d\})$ random vectors. In addition, let

$$(2.5) \quad W_j \equiv (w_{1,j}, \dots, w_{d,j})^T = (\lambda_1^{-\frac{1}{2}} z_{1,j}, \dots, \lambda_d^{-\frac{1}{2}} z_{d,j})^T, \quad j = 1, \dots, n,$$

and the W_j are iid $N(0, I_d)$ random vectors, where I_d is the d -dimensional identity matrix.

The following additional conditions are needed to ensure the consistency of \hat{u}_1^{ST} :

- The non-zero entries of the population eigenvector u_1 need to be a certain distance away from zero. In fact, if the non-zero entries of the first population eigenvector are close to zero, the corresponding entries of the first empirical eigenvector would also be small and look like pure noise entries. Thus, we assume

$$\max_{1 \leq i \leq \lfloor d^\beta \rfloor} |u_{i,1}|^{-1} \sim d^{\frac{\eta}{2}}, \quad \text{where } \eta \in [0, \alpha).$$

- From (2.4), we have

$$X_j = \sum_{i=1}^d z_{i,j} u_i, \quad j = 1, \dots, n.$$

Since $z_{1,j}$ has the largest variance λ_1 , then $z_{1,j} u_1$ contributes the most to the variance of X_j , $j = 1, \dots, n$. Note that $z_{1,j} u_1$ is consistent with u_1 , and so $z_{1,j} u_1$ is the key to making the simple thresholding method work. So we need to show that the remaining parts

$$(2.6) \quad H_j \equiv (h_{1,j}, \dots, h_{d,j})^T = \sum_{i=2}^d z_{i,j} u_i, \quad j = 1, \dots, n$$

have a negligible effect on the direction vector \hat{u}_1^{ST} .

- Suppose that the H_j are iid $N(0, \Delta_d)$, where $\Delta_d = (m_{kl})_{d \times d}$, for $j = 1, \dots, n$. A sufficient condition to make their effect negligible is the following mixing condition of Leadbetter, Lindgren and Rootzen (1983) [?]:

$$(2.7) \quad |m_{kl}| \leq m_{kk}^{\frac{1}{2}} m_{ll}^{\frac{1}{2}} \rho_{|k-l|}, \quad 1 \leq k \neq l \leq \lfloor d^\beta \rfloor,$$

where $\rho_t < 1$ for all $t > 1$ and $\rho_t \log(t) \rightarrow 0$, as $t \rightarrow \infty$. This mixing condition can guarantee that $\max_{1 \leq j \leq n} |h_{1,j}|$ has a quick convergence rate, as $d \rightarrow \infty$. It enables us to neglect the influence of H_j for sufficiently large d and make $z_{i,j} u_1$ the dominant component, which then gives consistency to the first population eigenvector u_1 . Thus the thresholding estimator \hat{u}_1^{ST} becomes consistent.

We now state one of the main theorems:

THEOREM 2.2. *Assume that X_1, \dots, X_n are random samples from a d -dimensional normal distribution $N(0, \Sigma_d)$. Define Z_j, W_j and H_j as in (2.4), (2.5), and (2.6) for $j = 1, \dots, n$. The first population eigenvector is $u_1 = (u_{1,1}, \dots, u_{d,1})^T$ with $u_{i,1} \neq 0, i = 1, \dots, \lfloor d^\beta \rfloor$, and otherwise $u_{i,1} = 0$.*

If the following conditions are satisfied:

- (a) $\lambda_1 \sim d^\alpha$, $\lambda_2 \sim d^\theta$, and $\sum_{i=2}^d \lambda_i \sim d$, where $\theta \in [0, \alpha)$ and $\alpha \in (0, 1]$,
- (b) for a non-negative constant C , $\lim_{d \rightarrow \infty} \frac{\lambda_1}{\sum_{i=2}^d \lambda_i} = C$ and ε -condition (2.3) is satisfied,
- (c) $\max_{1 \leq i \leq \lfloor d^\beta \rfloor} |u_{i,1}|^{-1} \sim d^{\frac{\eta}{2}}$, where $\eta \in [0, \alpha)$,
- (d) H_j satisfies the mixing condition (2.7), $j = 1, \dots, n$,
- (e) $\log^\delta(d) d^{\frac{\theta}{2}} \leq \lambda \leq d^{\frac{\gamma}{2}}$, where $\delta \in (\frac{1}{2}, \infty)$ and $\gamma \in (\theta, \alpha - \eta)$,

then the thresholding estimator \hat{u}_1^{ST} is consistent with u_1 .

We offer a couple of remarks regarding the above theorem. First of all, the theorem naturally reduces to Theorem 2.1 if we let the sparsity index $\beta = 0$. More importantly, this theorem, and the following ones in Sections 2 to 4, show that the concepts depicted in Figure 1 hold much more generally than just under the conditions of Example 1.1. In particular, in the above Theorem 2.2, setting $\theta = 0$ and $\eta = \beta$ would give the results plotted in Figure 1.

In addition, for different thresholding parameter λ , the ST estimator \hat{u}_1^{ST} is consistent with u_1 with different convergence rate. This result is stated in the following theorem. The notation $\lambda = o(d^\rho)$ below means that $\lambda d^{-\rho} \rightarrow 0$ as $d \rightarrow \infty$.

THEOREM 2.3. *For the thresholding parameter $\lambda = o(d^{\frac{\alpha-\eta-\varsigma}{2}})$, where $\varsigma \in [0, \alpha - \eta - \theta)$, the corresponding thresholding estimator \hat{u}_1^{ST} is consistent with u_1 , with a convergence rate of $d^{\frac{\varsigma}{2}}$.*

3. Asymptotic properties of RSPCA. As noted in Section 1, several sparse PCA methods have been proposed in the literature. Here we perform a detailed HDLSS asymptotic analysis of the sparse PCA procedure developed by Shen and Huang (2008) [?]. For simplicity, we refer to it as the *regularized sparse PCA*, or RSPCA for short. All the detailed technical proofs are again provided in Section 7 and the supplement materials.

We start with briefly reviewing the methodological details of RSPCA. (For more details, see [?].) Given a d -by- n data matrix $X_{(d)}$, consider the following penalized sum-of-squares criterion:

$$(3.1) \quad \|X_{(d)} - uv^T\|_F^2 + P_\lambda(u), \quad \text{subject to} \quad \|v\| = 1,$$

where u is a d -vector, v is a unit n -vector, $\|\cdot\|_F$ denotes the Frobenius norm, and $P_\lambda(u) = \sum_{i=1}^d p_\lambda(|u_{i,1}|)$ is a penalty function with $\lambda \geq 0$ being the penalty parameter. The penalty function can be any sparsity-inducing penalty. In particular, Shen and Huang [?] considered the soft thresholding (or L_1 or LASSO) penalty of Tibshirani (1996) [?], the hard thresholding penalty of Donoho and Johnstone (1994) [?], and the smoothly clipped absolute deviation (SCAD) penalty of Fan and Li (2001) [?].

Without the penalty term or when $\lambda = 0$, minimization of (3.1) can be obtained via singular value decomposition (SVD) [?], which results in the best rank-one approximation of $X_{(d)}$ as $\tilde{u}_1 \tilde{v}_1^T$, where \tilde{u}_1 and \tilde{v}_1^T minimize the criterion (3.1). The normalized \tilde{u}_1 turns out to be the first empirical PC loading vector. With the penalty term, Shen and Huang define the sparse PC loading vector as $\hat{u}_1 = \tilde{u}_1 / \|\tilde{u}_1\|$ where \tilde{u}_1 is now the minimizer of (3.1) with the penalty term included. The minimization now needs to be performed iteratively. For a given \tilde{v}_1 in the criterion (3.1), we can get the minimizing vector as $\tilde{u}_1 = h_\lambda(X_{(d)}\tilde{v}_1)$, where h_λ is a thresholding function that depends on the particular penalty function used and the penalty (or thresholding) parameter λ . See [?] for more details. The thresholding is applied to the vector $X_{(d)}\tilde{v}_1$ componentwise.

Shen and Huang (2008) [?] proposed the following iterative procedure for minimizing the criterion (3.1):

The RSPCA Algorithm

-
1. Initialize:

- (a) Use SVD to obtain the best rank-one approximation $\tilde{u}_1 \tilde{v}_1^T$ of the data matrix $X_{(d)}$, where \tilde{v}_1 is a unit vector.
- (b) Set $\tilde{u}_1^{\text{old}} = \tilde{u}_1$ and $\tilde{v}_1^{\text{old}} = \tilde{v}_1$.
2. Update:
 - (a) $\tilde{u}_1^{\text{new}} = h_\lambda(X_{(d)} \tilde{v}_1^{\text{old}})$.
 - (b) $\tilde{v}_1^{\text{new}} = \frac{X_{(d)}^T \tilde{u}_1^{\text{new}}}{\|X_{(d)}^T \tilde{u}_1^{\text{new}}\|}$.
3. Repeat Step 2 setting $\tilde{u}_1^{\text{old}} = \tilde{u}_1^{\text{new}}$ and $\tilde{v}_1^{\text{old}} = \tilde{v}_1^{\text{new}}$ until convergence.
4. Normalize the final \tilde{u}_1^{new} to get \hat{u}_1 , the desired sparse loading vector.

There exists a nice connection between the simple thresholding (ST) method of Section 2 and RSPCA. The ST estimator \hat{u}_1^{ST} is exactly the sparse loading vector \hat{u}_1 obtained from the first iteration of the RSPCA iterative algorithm, when the hard thresholding penalty is used. In particular, the first dual eigenvector \tilde{v}_1 in (2.1) is just the \tilde{v}_1 from the best rank-one approximation $\tilde{u}_1 \tilde{v}_1^T$ of the data matrix $X_{(d)}$. Then, the application of the simple thresholding method to the vector $X_{(d)} \tilde{v}_1$ as in (2.2) leads to the sparse ST estimator \hat{u}_1^{ST} . This is the same as applying the hard thresholding penalty in (3.1) to generate the sparse loading vector \hat{u}_1 , for the given \tilde{v}_1 . The thresholding parameter λ in (2.2) also corresponds to the penalty parameter λ in (3.1) in the case of the hard thresholding penalty.

Below we develop conditions under which the sparse RSPCA loading vector \hat{u}_1 is consistent with the population eigenvector u_1 when a proper thresholding parameter λ is used. All three of the soft thresholding, hard thresholding or SCAD penalties are considered. First, the following theorem states conditions when the first step sparse loading vector \hat{u}_1 is consistent with u_1 under the proper thresholding parameter λ .

THEOREM 3.1. *Under the assumptions and conditions of Theorem 2.2, the first step sparse loading vector \hat{u}_1 is consistent with u_1 .*

Theorem 3.1 explores conditions when the first iteration of the iterative procedure of RSPCA gives a consistent sparse loading vector \hat{u}_1 , with an appropriate thresholding parameter λ . Similar to the ST estimator, for different parameters λ , \hat{u}_1 is consistent with u_1 with different convergence rates. The result is given in the following Theorem 3.2.

THEOREM 3.2. *For the thresholding parameter $\lambda = o(d^{\frac{\alpha-\eta-\varsigma}{2}})$, where $\varsigma \in [0, \alpha - \eta - \theta)$, the sparse loading vector \hat{u}_1 in Theorem 3.1 is consistent with u_1 , with a convergence rate of $d^{\frac{\varsigma}{2}}$.*

We then set \hat{u}_1^{old} to be the consistent sparse loading vector obtained after the first iteration of the RSPCA algorithm. We then obtain an updated estimate for v_1 as $\tilde{v}_1^{\text{new}} = X_{(d)}^T \hat{u}_1^{\text{old}} / \|X_{(d)}^T \hat{u}_1^{\text{old}}\|$. The theorem below studies the asymptotic properties of \tilde{v}_1^{new} .

THEOREM 3.3. *Assume that \hat{u}_1^{old} is consistent with u_1 with the convergence rate $d^{\frac{\varsigma}{2}}$, where $\varsigma \in [1 - \alpha, \infty)$. If the ε -condition is satisfied, then*

$$\tilde{v}_1^{\text{new}} \xrightarrow{p} \frac{\tilde{W}_1}{\|\tilde{W}_1\|}, \quad \text{as } d \rightarrow \infty,$$

where $\tilde{W}_1 = (w_{1,1}, \dots, w_{1,n})$ follows a standard n -dimensional normal distribution $N(0, I_n)$ and the $w_{i,j}$ are defined in (2.5).

Since Theorem 3.3 establishes the asymptotic properties of \tilde{v}_1^{new} , we can now study the asymptotic properties of the updated sparse loading vector

$$(3.2) \quad \hat{u}_1^{\text{new}} = \frac{\tilde{u}_1^{\text{new}}}{\|\tilde{u}_1^{\text{new}}\|}, \quad \text{with } \tilde{u}_1^{\text{new}} = h_\lambda(X_{(d)} \tilde{v}_1^{\text{new}}),$$

as defined in the iterative procedure of RSPCA. The following Theorem 3.4 shows that with a proper choice of the thresholding parameter λ , the updated sparse loading vector \hat{u}_1^{new} remains to be consistent with the population eigenvector u_1 .

THEOREM 3.4. *Under the assumptions and conditions of Theorems 2.2 and 3.3, the updated sparse loading vector \hat{u}_1^{new} in (3.2) is consistent with u_1 .*

For different threshold parameters λ , \hat{u}_1^{new} is again consistent with u_1 with different convergence rates, as seen in the following theorem.

THEOREM 3.5. *For the thresholding parameter $\lambda = o(d^{\frac{\alpha - \eta - \varsigma}{2}})$, where $\varsigma \in [0, \alpha - \eta - \theta)$, the updated sparse loading vector \hat{u}_1^{new} in Theorem 3.4 is consistent with u_1 , with convergence rate $d^{\frac{\varsigma}{2}}$.*

According to Theorems 3.2 and 3.5, if $\alpha - \eta - \theta > 1 - \alpha$, then we can choose the thresholding parameter $\lambda = o(d^{\frac{\alpha - \eta - \varsigma}{2}})$ and make the updated sparse loading vector \hat{u}_1^{new} in (3.2) to be consistent with u_1 at every updating step.

4. Strong Inconsistency. We have shown that we can attain consistency using sparse PCA, when the spike index α is greater than the sparsity index β . This motivates the question of consistency using sparse PCA when the spike index α is smaller than the sparsity index β . To answer this question, we consider an oracle estimator which uses the exact positions of zero entries of the population eigenvector u_1 . We will show that even this oracle estimator is strongly inconsistent with the population eigenvector u_1 when the spike index α is smaller than the sparsity index β . Compared with this oracle sparse PCA, threshold methods can perform no better because they also need to estimate location of the zero entries; hence threshold methods will also be strongly inconsistent.

For Example 1.1, the first $\lfloor d^\beta \rfloor$ entries of the population eigenvector u_1 are known to be the non-zero entries. So we could first find a $\lfloor d^\beta \rfloor$ -dimensional estimator \hat{u}_1^* through subspace PCA for the $\lfloor d^\beta \rfloor$ -dimensional subspace eigenvector u_1^* which is proportional to the following $\lfloor d^\beta \rfloor$ -dimensional vector,

$$\ddot{u}_1 = (\overbrace{1, \dots, 1}^{\lfloor d^\beta \rfloor})^T.$$

Then we get the oracle (OR) estimator for u_1 as,

$$\hat{u}_1^{\text{OR}} = ((\hat{u}_1^*)^T, \overbrace{0, \dots, 0}^{d - \lfloor d^\beta \rfloor})^T.$$

The oracle estimator \hat{u}_1^{OR} has the same sparsity as the population eigenvector u_1 . Furthermore, it is strongly inconsistent with u_1 when $\alpha < \beta$.

To make this precise, we study the procedure to generate the oracle estimator for general single component models. Assume that the first $\lfloor d^\beta \rfloor$ entries of the population eigenvector u_1 are non-zero and the rest are all zero: $u_1 = (u_{1,1}, \dots, u_{d,1})^T$, where $u_{i,1} \neq 0, i = 1, \dots, \lfloor d^\beta \rfloor$, otherwise $u_{i,1} = 0$. Let $X_j^* = (x_{1,j}, \dots, x_{\lfloor d^\beta \rfloor, j})^T \sim N(0, \Sigma_{\lfloor d^\beta \rfloor}^*)$, where $\Sigma_{\lfloor d^\beta \rfloor}^*$ is the covariance matrix of $X_j^*, j = 1, \dots, n$. Then, the eigen-decomposition of $\Sigma_{\lfloor d^\beta \rfloor}^*$ is

$$\Sigma_{\lfloor d^\beta \rfloor}^* = U_{\lfloor d^\beta \rfloor}^* \Lambda_{\lfloor d^\beta \rfloor}^* (U_{\lfloor d^\beta \rfloor}^*)^T,$$

where Λ_d^* is the diagonal matrix of eigenvalues $\lambda_1^* \geq \lambda_2^* \geq \dots \geq \lambda_{\lfloor d^\beta \rfloor}^*$ and $U_{\lfloor d^\beta \rfloor}^*$ is the matrix of the corresponding eigenvectors so that $U_{\lfloor d^\beta \rfloor}^* = [u_1^*, \dots, u_{\lfloor d^\beta \rfloor}^*]$. Since the last $d - \lfloor d^\beta \rfloor$ entries of the first population eigenvector u_1 equal zero, it follows that the first eigenvector u_1^* of $\Sigma_{\lfloor d^\beta \rfloor}^*$ is

formed by the non-zero entries of the population eigenvector u_1 , i.e. $u_1^* = (u_{1,1}, \dots, u_{\lfloor d^\beta \rfloor, 1})^T$. So we have

$$u_1 = ((u_1^*)^T, \overbrace{0, \dots, 0}^{d - \lfloor d^\beta \rfloor})^T.$$

Consider the following data matrix $X_{\lfloor d^\beta \rfloor}^* = [X_1^*, \dots, X_n^*]$, and denote the sample covariance matrix by $\hat{\Sigma}_{\lfloor d^\beta \rfloor}^* = n^{-1} X_{\lfloor d^\beta \rfloor}^* X_{\lfloor d^\beta \rfloor}^{*T}$. Then, the sample covariance matrix $\hat{\Sigma}_{\lfloor d^\beta \rfloor}^*$ can be similarly decomposed as

$$\hat{\Sigma}_{\lfloor d^\beta \rfloor}^* = \hat{U}_{\lfloor d^\beta \rfloor}^* \hat{\Lambda}_{\lfloor d^\beta \rfloor}^* (\hat{U}_{\lfloor d^\beta \rfloor}^*)^T,$$

where $\hat{\Lambda}_{\lfloor d^\beta \rfloor}^*$ is the diagonal matrix of the sample eigenvalues $\hat{\lambda}_1^* \geq \hat{\lambda}_2^* \geq \dots \geq \hat{\lambda}_{\lfloor d^\beta \rfloor}^*$ and $\hat{U}_{\lfloor d^\beta \rfloor}^*$ is the matrix of the corresponding sample eigenvectors so that $\hat{U}_{\lfloor d^\beta \rfloor}^* = [\hat{u}_1^*, \dots, \hat{u}_d^*]$. Then, we define the oracle (OR) estimator as

$$(4.1) \quad \hat{u}_1^{\text{OR}} = ((\hat{u}_1^*)^T, \overbrace{0, \dots, 0}^{d - \lfloor d^\beta \rfloor})^T.$$

The following theorem states the main result regarding strong inconsistency.

THEOREM 4.1. *Assume that X_1, \dots, X_n are random samples from a d -dimensional normal distribution $N(0, \Sigma_d)$. The first population eigenvector is $u_1 = (u_{1,1}, \dots, u_{d,1})^T$, where $u_{i,1} \neq 0, i = 1, \dots, \lfloor d^\beta \rfloor$, otherwise $u_{i,1} = 0$. If the following conditions are satisfied:*

- (a) $\lambda_1 \sim d^\alpha$, $\lambda_2 \sim d^\theta$, $\lambda_d \sim 1$ and $\sum_{i=2}^d \lambda_i \sim d$, where $\theta \in [0, \frac{\beta}{2})$,
- (b) $\alpha < \beta$,

then the oracle estimator \hat{u}_1^{OR} in (4.1) is strongly inconsistent with u_1 .

5. Simulations for sparse PCA. Here, we will perform simulation studies to illustrate the performance of the ST method and the RSPCA with the hard thresholding penalty. Let the sample size $n = 25$ and the dimension $d = 10,000$. To generate the data matrix $X_{(d)}$, we first need to construct the population covariance matrix for $X_{(d)}$ that approximates the conditions of Theorems 2.2 and 3.1 when the spike index α is greater than the sparsity index β .

For the population covariance matrix, we consider the motivating model in Example 1.1 for the first population eigenvector and the eigenvalues, where the first eigenvalue $\lambda_1 = d^\alpha$ and the rest equal one, i.e. $\lambda_i = 1, i \geq 2$. For

the additional population eigenvectors u_i , $2 \leq i \leq \lfloor d^\beta \rfloor$, let the last $d - \lfloor d^\beta \rfloor$ entries of these eigenvectors be zero. In particular, let the eigenvectors u_i , $2 \leq i \leq \lfloor d^\beta \rfloor$, be proportional to

$$\dot{u}_i = (\overbrace{1, \dots, 1}^{i-1}, -i+1, 0, \dots, 0)^T.$$

After normalizing \dot{u}_i , we get the i -th eigenvector $u_i = \dot{u}_i / \|\dot{u}_i\|$. For $i > \lfloor d^\beta \rfloor$, let the i -th eigenvector have just one non-zero entry in the i -th position such

that $u_i = (\overbrace{0, \dots, 0}^{i-1}, 1, 0, \dots, 0)^T$.

Then the data matrix is generated as

$$X_{(d)} = d^{\frac{\alpha}{2}} u_1 z_1^T + \sum_{i=2}^d u_i z_i^T,$$

where the z_i are generated from the n -dimensional standard normal distribution $N(0, I_n)$.

We select twenty spike and sparsity pairs (α, β) that have spike index $\alpha = \{0.2, 0.4, 0.6, 0.8\}$ and sparsity index $\beta = \{0, 0.1, 0.3, 0.5, 0.7\}$, which are shown in Figure 1. We perform the simulation for all twenty spike and sparsity pairs. For each spike and sparsity pair (α, β) , we generate 100 realizations of the data matrix $X_{(d)}$. Results for three representative pairs are reported below, and interesting observations are discussed. Additional simulation results can be found at [?].

First of all, the plots in Figure 2 summarize the results for the spike and sparsity pair $(\alpha, \beta) = (0.6, 0.1)$, corresponding to one of the square dots in the white (consistent) triangular area of Figure 1. For each replication of the data matrix $X_{(d)}$ and a range of the thresholding parameter λ , we obtained the ST estimator \hat{u}_1^{ST} (Section 2) and the RSPCA estimator \hat{u}_1 (Section 4). Then we calculate the angle between the estimates \hat{u}_1^{ST} (or \hat{u}_1) and the first population eigenvector u_1 through (1.1). Plotting this angle as a function of the thresholding parameter λ gives the curve in Panel (A) of Figure 2. Since ST and RSPCA have very similar performance in this case, we just show the RSPCA plots in Figure 2. The 100 simulation realizations of the data matrices $X_{(d)}$ generate the one-hundred curves in the panel. We rescale the thresholding parameter λ as $\log_{10}(\lambda + 10^{-5})$, to help reveal clearly the tendency of the angle curves as the thresholding parameter increases.

In these angle plots, the angles with $\lambda = 0$ (essentially the left edge of each plot) correspond to the ones obtained by the conventional PCA. Note




FIG 2. Performance summary of RSPCA for spike index $\alpha = 0.6$ and sparsity index $\beta = 0.1$ where consistency is expected. Panel (A) shows angle to the first population eigenvector as a function of thresholding parameter λ . Panel (B) and (C) are Type I Error and Type II Error as a function of λ . The vertical dashed and solid lines are the left and right bounds of the range of the thresholding parameter, which leads to the consistency of RSPCA. These show very good performance of RSPCA within the indicated range, which empirically confirms our asymptotic calculation. The circles indicate values at the BIC choice of λ .

that these angles are all over 40 degrees which confirms the results of Jung and Marron (2009) [?] that when the spike index $\alpha < 1$, the conventional PCA can not generate a consistent estimator for the population eigenvector u_1 . As λ increases, the angle remains stable for a while, then decreases to almost 0 degree, before eventually starting to increase to 90 degrees. The dashed and solid vertical lines in the angle plots indicate the range of the thresholding parameter that gives a consistent estimator for u_1 , as stated in Theorems 2.2 and 3.4. These plots suggest that RSPCA does improve over PCA and the indicated thresholding range is very reasonable in this case, which in turn empirically validates the asymptotic results of the theorems. For each realization of the data, as the thresholding parameter increases, all entries will be thresholded out, i.e. become zero, so the sparse PCA estimator eventually becomes a d -dimensional zero vector. Hence the angles go to 90 degrees when the thresholding parameter is large enough.

Zou, Hastie and Tibshirani (2007) [?] suggest the use of the Bayesian Information Criterion (BIC) [?] to select the number of the non-zero coefficients for a lasso regression. Lee et al. (2010) [?] apply this idea to the sparse PCA context. According to [?], for a fixed \tilde{v}_1 , minimization of (3.1) with respect to \tilde{u}_1 is equivalent to minimization of the following penalized regression criterion with respect to \tilde{u}_1 :

$$(5.1) \quad \|X_{(d)} - \tilde{u}_1 \tilde{v}_1^T\|_F^2 + P_\lambda(\tilde{u}_1) = \|Y - (I_d \otimes \tilde{v}_1) \tilde{u}_1\|^2 + P_\lambda(\tilde{u}_1),$$

where $Y = (X^1, \dots, X^d)^T$, with X^i being the i -th row of $X_{(d)}$, and \otimes is the Kronecker product. Following their suggestion, for the above penalized regression (5.1) with a fixed \tilde{v}_1 , we define

$$(5.2) \quad \text{BIC}(\lambda) = \frac{\|Y - \hat{Y}\|^2}{nd\hat{\sigma}^2} + \frac{\log(nd)}{nd} \hat{d}f(\lambda),$$

where $\hat{\sigma}^2$ is the ordinary-least squares estimate of the error variance, and $\hat{d}f(\lambda)$ is the degree of sparsity for the thresholding parameter λ , i.e. the number of non-zero entries in \tilde{u}_1 . For every step of the iterative procedure of RSPCA, we can use BIC (5.2) to select the thresholding parameter and then obtain the corresponding sparse PC direction, until the algorithm converges.

For every angle curve in the angle plots of Figure 2, we use a blue circle to indicate the thresholding parameter λ that is selected by BIC during the last iterative step of RSPCA, and the corresponding angle. In the current $\alpha = 0.6$, $\beta = 0.1$ context, BIC works well, and all the BIC-selected λ values are very close, so the 100 circles are essentially over plotted on each other. BIC also works well for the other spike and sparsity pairs (α, β) we considered where $\alpha > \beta$, which are shown in [?].

Another measure of the success of a sparse estimator is in terms of which entries are zeroed. Type I Error is the proportion of non-zero entries in u_1 that are mistakenly estimated as zero. Type II Error is the proportion of zero entries in u_1 that are mistakenly estimated as non-zero. Similar to the angle, Type I Error (Type II Error) is also a function of the thresholding parameter. For each replication of the data matrix $X_{(d)}$, we calculate a Type I Error (Type II Error) curve. Thus, there are one hundred such curves in Panels (B) and (C) of Figure 2, respectively. The dashed and solid lines in these two panels are the same as those in Panel (A). Note that for all the thresholding parameters in the range indicated by the lines, the errors are very small, which is again consistent with the asymptotic results of Theorems 2.2 and 3.4. Again, the circles in these plots are selected by BIC and they have the same horizontal thresholding parameter, as in the angle plots. Thus, BIC works well here. BIC also generates similarly very small errors for the other spike and sparsity pairs (α, β) in Figure 1 that satisfy $\alpha > \beta$.

Next we will compare the relative performance among PCA, ST and RSPCA. In almost all cases, ST and RSPCA give better results than PCA and in some extreme cases, the three methods have similar poor performance. Although in most cases both ST and RSPCA have similar performance, however, there are some cases (for example when $\alpha = 0.4$ and $\beta = 0.3$), where RSPCA performs better than ST. For every replication of

the data matrix $X_{(d)}$, we use BIC to select the thresholding parameter, and then calculate the ST estimator \hat{u}_1^{ST} and the RSPCA estimator \hat{u}_1 . After that, we calculate the angle, Type I Error and Type II Error for the three estimators, as well as the difference between ST and RSPCA (ST minus RSPCA). For each measure, the 100 values are summarized using box plots in Figure 3.




FIG 3. Comparison of PCA, ST and RSPCA for spike index $\alpha = 0.4$ and sparsity index $\beta = 0.3$. Panels (A), (B) and (C) respectively contain four angle, Type I Error and Type II Error box plots: (i) conventional PCA; (ii) and (iii) ST and RSPCA with BIC; (iv) the difference between ST and RSPCA. In Panel (A), angles for conventional PCA are generally larger than ST and RSPCA which indicates the worse performance of conventional PCA. In addition, the angles and Type I Errors for ST are larger than RSPCA and their difference box plots furthermore confirm this point, which indicates the better performance of RSPCA in this case. Type II Errors for ST and RSPCA are almost the same.

Panel (A) of Figure 3 shows the box plots of the angles between the first population eigenvector u_1 and the estimates obtained by PCA, ST and RSPCA, as well as the differences between ST and RSPCA. Note that the PCA angles are large, compared with ST and RSPCA, indicating the worse performance of PCA. The angle of ST seems larger than RSPCA. For a deeper view of this comparison, the pairwise differences are studied in the fourth box plot of the panel. The angle differences are almost always positive, with some differences bigger than 50 degrees, which suggests that RSPCA has a better performance than ST. Similar conclusions can be made from the box plots of the errors, in Panels (B) and (C) of Figure 3. The box plot for PCA is not shown in Panel (C) because the corresponding Type II Error almost always equals one, which is far outside the shown range of interest.

Finally, Theorems 2.2 and 3.4 consider the condition that the spike index α is greater than the sparsity index β . When α is smaller than β , neither ST nor RSPCA is expected to give consistent estimation for the first population




FIG 4. Performance summary of RSPCA for spike index $\alpha = 0.2$ and sparsity index $\beta = 0.7$ where strong inconsistency is expected. Panel (A) shows the angle curves, between the RSPCA estimator and the first population eigenvector. Same format as Figure 2. Since the spike index $\alpha = 0.2$ is smaller than the sparsity index $\beta = 0.7$, it follows that the right bound (solid line) is smaller than the left bound (dashed line). Thus the theorems do not give a meaningful range of the thresholding parameter. As expected, performance is very poor for any thresholding parameter λ .

eigenvector u_1 , as discussed in Section 4. For the spike and sparsity pairs (α, β) such that $\alpha < \beta$, the simulation results also confirm this point. Here, we display the simulation plots for the spike and sparsity pair $(\alpha, \beta) = (0.2, 0.7)$ in Figure 4 as a representative of such simulations. Since ST and RSPCA have very similar performance here, we just show the simulation results for RSPCA. Similar to Figure 2, the circles in Figure 4 correspond to the thresholding parameter selected by BIC. From the angle plots, we can see that the angles, selected by BIC, are close to 90 degrees, which suggests the failure of BIC in this case. In fact, all the angle curves are above 80 degrees. Thus, neither ST nor RSPCA generates a reasonable sparse estimator. This is a common phenomenon when the spike index α is smaller than the sparsity index β . It is consistent with the theoretical investigation in Section 4.

Furthermore, the corresponding Type I Error, generated by ST or RSPCA with BIC, is close to one. This further confirms that BIC doesn't work when the spike index α is smaller than the sparsity index β . ST and RSPCA with $\lambda = 0$ is just the conventional PCA, and typically will not generate a sparse estimator. This entails that the Type I Error and Type II Error, corresponding to $\lambda = 0$, respectively equals zero and one. As the thresholding parameter increases, more and more entries are thresholded out; hence Type I Error increases to one and Type II Error decreases to zero.

6. Non-Gaussian Variations. In this paper, we consider HDLSS data analysis contexts, using the high dimensional normal distribution. In the fu-

ture, we hope to extend our theorems to more general distributions. However, this will be challenging because sparse PCA methods may not work in some extreme cases. This point is illustrated by the following interesting example.

EXAMPLE 6.1. *Let $\alpha \in (0, 1)$ and $X = (x_1, \dots, x_d)^T$, where $\{x_i, i = 1, \dots, d\}$ are independent discrete random variables with the following discrete probability distributions:*

$$x_1 = \begin{cases} d^{\frac{\alpha}{2}}, & \text{with probability } \frac{1}{2}, \\ -d^{\frac{\alpha}{2}}, & \text{with probability } \frac{1}{2}; \end{cases}$$

and for $i = 2, \dots, d$

$$x_i = \begin{cases} d^{\frac{\alpha+1}{4}}, & \text{with probability } d^{-\frac{\alpha+1}{2}}, \\ -d^{\frac{\alpha+1}{4}}, & \text{with probability } d^{-\frac{\alpha+1}{2}}, \\ 0, & \text{with probability } 1 - 2d^{-\frac{\alpha+1}{2}}. \end{cases}$$

Then X has mean 0 and variance-covariance Σ_d with

$$\Sigma_d = d^\alpha u_1 u_1^T + \sum_{k=2}^d u_k u_k^T,$$

where $u_1 = (1, 0, \dots, 0)^T$.

Suppose that we only have sample size $n = 1$, i.e. $X_1 = (x_{i,1}, \dots, x_{d,1})^T$, then the first empirical eigenvector

$$\hat{u}_1 = (\hat{u}_{1,1}, \dots, \hat{u}_{d,1})^T = \frac{1}{\sqrt{\sum_{i=1}^d x_{i,1}^2}} (x_{i,1}, \dots, x_{i,d})^T.$$

Under this condition, we have

$$\begin{aligned} P(\arg\max_i |\hat{u}_{i,d}| = 1) &= P(|x_{1,1}| > \max\{|x_{2,1}|, \dots, |x_{d,1}|\}) \\ &= P(x_{2,1} = 0, \dots, x_{d,1} = 0) \\ &= \left(1 - 2d^{-\frac{\alpha+1}{2}}\right)^{d-1} \\ &\longrightarrow 0 \text{ as } d \rightarrow \infty. \end{aligned}$$

In particular, the absolute value of the first entry of the empirical eigenvector can not be greater than the others with probability 1, so we can not always threshold out the right entries which results in the failure of the simple thresholding method. Similar considerations apply to other sparse PCA methods.

7. Proofs.

7.1. *Proofs of Theorem 2.2 and Theorem 2.3.* In order to prove Theorem 2.2 and Theorem 2.3, we need the dependent extreme value results from Leadbetter, Lindgren and Rootzen (1983) [?], in particular their Lemma 6.1.1 and Theorem 6.1.3.

An immediate consequence of those results is the following proposition.

PROPOSITION 7.1. *Suppose that the standard normal sequence $\{\xi_i, i = 1, \dots, \lfloor d^\beta \rfloor\}$ satisfies the mixing condition (2.7). Let the positive constants $\{c_i\}$ be such that $\sum_{i=1}^{\lfloor d^\beta \rfloor} (1 - \Phi(c_i))$ is bounded and such that $C_{\lfloor d^\beta \rfloor} = \min_{1 \leq i \leq \lfloor d^\beta \rfloor} c_i \geq c(\log(\lfloor d^\beta \rfloor))^{\frac{1}{2}}$ for some $c > 0$.*

Then the following holds:

$$P \left[\bigcap_{i=1}^{\lfloor d^\beta \rfloor} \{\xi_i \leq c_i\} \right] - \prod_{i=1}^{\lfloor d^\beta \rfloor} \Phi(c_i) \longrightarrow 0, \text{ as } d \rightarrow \infty,$$

where Φ is the standard normal distribution function. Furthermore, if for some $j \geq 0$, we have

$$\sum_{i=1}^{\lfloor d^\beta \rfloor} (1 - \Phi(c_i)) \longrightarrow j, \text{ as } d \rightarrow \infty,$$

then

$$P \left[\bigcap_{i=1}^{\lfloor d^\beta \rfloor} \{\xi_i \leq c_i\} \right] \longrightarrow e^{-j}, \text{ as } d \rightarrow \infty.$$

Proposition 7.1 is used to control the right side of (2.6) through the following lemma.

LEMMA 7.1. *Suppose that $\xi_i \sim N(0, \delta_{i,i})$ satisfies the mixing condition (2.7), where δ_{ij} is the covariance of the normal sequence $\{\xi_i\}$, $i, j = 1, \dots, \lfloor d^\beta \rfloor$. If $C_{\lfloor d^\beta \rfloor} \geq (\log(\lfloor d^\beta \rfloor))^\delta \max_{1 \leq i \leq \lfloor d^\beta \rfloor} \delta_{ii}^{\frac{1}{2}}$, where $\delta \in (\frac{1}{2}, \infty)$, then*

$$C_{\lfloor d^\beta \rfloor}^{-1} \max_{1 \leq i \leq \lfloor d^\beta \rfloor} |\xi_i| \xrightarrow{p} 0, \text{ as } d \rightarrow \infty.$$

PROOF. Note that for every $\tau > 0$

$$\begin{aligned}
 (7.1) \quad & P \left[C_{\lfloor d^\beta \rfloor}^{-1} \max_{1 \leq i \leq \lfloor d^\beta \rfloor} |\xi_i| > \tau \right] = P \left[\max_{1 \leq i \leq \lfloor d^\beta \rfloor} |\xi_i| > C_{\lfloor d^\beta \rfloor} \tau \right] \\
 & \leq P \left[\left\{ \max_{1 \leq i \leq \lfloor d^\beta \rfloor} \xi_i > C_{\lfloor d^\beta \rfloor} \tau \right\} \cup \left\{ \max_{1 \leq i \leq \lfloor d^\beta \rfloor} (-\xi_i) > C_{\lfloor d^\beta \rfloor} \tau \right\} \right] \\
 & \leq P \left[\max_{1 \leq i \leq \lfloor d^\beta \rfloor} \xi_i > C_{\lfloor d^\beta \rfloor} \tau \right] + P \left[\max_{1 \leq i \leq \lfloor d^\beta \rfloor} (-\xi_i) > C_{\lfloor d^\beta \rfloor} \tau \right] \\
 & = 2P \left[\max_{1 \leq i \leq \lfloor d^\beta \rfloor} \xi_i > C_{\lfloor d^\beta \rfloor} \tau \right] \\
 & \leq 2 \left(1 - P \left[\bigcap_{i=1}^{\lfloor d^\beta \rfloor} \left\{ \xi_i \delta_{ii}^{-\frac{1}{2}} \leq c(\log(\lfloor d^\beta \rfloor))^\delta \right\} \right] \right),
 \end{aligned}$$

where c is a positive constant. Since

$$\sum_{i=1}^{\lfloor d^\beta \rfloor} \left(1 - \Phi \left(c(\log(\lfloor d^\beta \rfloor))^\delta \right) \right) \longrightarrow 0, \text{ as } d \rightarrow \infty,$$

it then follows from Proposition 7.1 that

$$(7.2) \quad P \left[\bigcap_{i=1}^{\lfloor d^\beta \rfloor} \left\{ \xi_i \delta_{ii}^{-\frac{1}{2}} \leq c(\log(\lfloor d^\beta \rfloor))^\delta \right\} \right] \longrightarrow 1, \text{ as } d \rightarrow \infty.$$

From (7.1) and (7.2), we can get

$$C_{\lfloor d^\beta \rfloor}^{-1} \max_{1 \leq i \leq \lfloor d^\beta \rfloor} |\xi_i| \xrightarrow{P} 0, \text{ as } d \rightarrow \infty.$$

□

Now we will begin the proof of Theorem 2.2 and Theorem 2.3. Denote $\tilde{X}_i = (x_{i,1}, \dots, x_{i,n})^T$, $\tilde{Z}_i = (z_{i,1}, \dots, z_{i,n})^T$, $\tilde{W}_i = (w_{i,1}, \dots, w_{i,n})^T$ and $\tilde{H}_i = (h_{i,1}, \dots, h_{i,n})$, $i = 1, \dots, d$.

Note that

$$(7.3) \quad | \langle \hat{u}_1^{\text{ST}}, u_1 \rangle | = \frac{|\sum_{i=1}^{\lfloor d^\beta \rfloor} \check{u}_{i,1} u_{i,1}|}{\sqrt{\sum_{i=1}^d (\check{u}_{i,1})^2}} = \frac{\lambda_1^{-\frac{1}{2}} |\sum_{i=1}^{\lfloor d^\beta \rfloor} \check{u}_{i,1} u_{i,1}|}{\lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^d (\check{u}_{i,1})^2}}.$$

Below we need to bound the denominator and the numerator of (7.3).

We start with the numerator. Since $\tilde{X}_i = u_{i,1}\tilde{Z}_1 + \tilde{H}_i, i = 1, \dots, d$, it follows that $\tilde{v}_1^T \tilde{X}_i = u_{i,1}\tilde{v}_1^T \tilde{Z}_1 + \tilde{v}_1^T \tilde{H}_i$, which yields

$$\begin{aligned}\check{u}_{i,1} &= u_{i,1}\tilde{v}_1^T \tilde{Z}_1 1_{\{|\tilde{v}_1^T \tilde{X}_i| > \lambda\}} + \tilde{v}_1^T \tilde{H}_i 1_{\{|\tilde{v}_1^T \tilde{X}_i| > \lambda\}} \\ &= u_{i,1}\tilde{v}_1^T \tilde{Z}_1 + u_{i,1}\tilde{v}_1^T \tilde{Z}_1 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}} + \tilde{v}_1^T \tilde{H}_i 1_{\{|\tilde{v}_1^T \tilde{X}_i| > \lambda\}},\end{aligned}$$

and

$$\begin{aligned}\sum_{i=1}^{[d^\beta]} \check{u}_{i,1} u_{i,1} &= \sum_{i=1}^{[d^\beta]} u_{i,1}^2 \tilde{v}_1^T \tilde{Z}_1 1_{\{|\tilde{v}_1^T \tilde{X}_i| > \lambda\}} + \sum_{i=1}^{[d^\beta]} u_{i,1} \tilde{v}_1^T \tilde{H}_i 1_{\{|\tilde{v}_1^T \tilde{X}_i| > \lambda\}} \\ &= \tilde{v}_1^T \tilde{Z}_1 + \sum_{i=1}^{[d^\beta]} u_{i,1}^2 \tilde{v}_1^T \tilde{Z}_1 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}} + \sum_{i=1}^{[d^\beta]} u_{i,1} \tilde{v}_1^T \tilde{H}_i 1_{\{|\tilde{v}_1^T \tilde{X}_i| > \lambda\}}.\end{aligned}$$

It follows that

$$\begin{aligned}(7.4) \quad \lambda_1^{-\frac{1}{2}} \left| \sum_{i=1}^{[d^\beta]} \check{u}_{i,1} u_{i,1} \right| &\leq \lambda_1^{-\frac{1}{2}} \sum_{i=1}^{[d^\beta]} u_{i,1}^2 |\tilde{v}_1^T \tilde{Z}_1| + \lambda_1^{-\frac{1}{2}} \sum_{i=1}^{[d^\beta]} |u_{i,1} \tilde{v}_1^T \tilde{H}_i| \\ &= |\tilde{v}_1^T \tilde{W}_1| + \sum_{i=1}^{[d^\beta]} \sum_{j=1}^n \lambda_1^{-\frac{1}{2}} |u_{i,1} h_{i,j}|,\end{aligned}$$

and

$$\begin{aligned}(7.5) \quad \lambda_1^{-\frac{1}{2}} \left| \sum_{i=1}^{[d^\beta]} \check{u}_{i,1} u_{i,1} \right| &\geq \lambda_1^{-\frac{1}{2}} |\tilde{v}_1^T \tilde{Z}_1| - \lambda_1^{-\frac{1}{2}} \sum_{i=1}^{[d^\beta]} u_{i,1}^2 |\tilde{v}_1^T \tilde{Z}_1| 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}} - \lambda_1^{-\frac{1}{2}} \sum_{i=1}^{[d^\beta]} |u_{i,1} \tilde{v}_1^T \tilde{H}_i| \\ &\geq |\tilde{v}_1^T \tilde{W}_1| - |\tilde{v}_1^T \tilde{W}_1| \sum_{i=1}^{[d^\beta]} u_{i,1}^2 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}} - \sum_{i=1}^{[d^\beta]} \sum_{j=1}^n \lambda_1^{-\frac{1}{2}} |u_{i,1} h_{i,j}|.\end{aligned}$$

Next we will show that

$$(7.6) \quad \sum_{i=1}^{[d^\beta]} \sum_{j=1}^n \lambda_1^{-\frac{1}{2}} |u_{i,1} h_{i,j}| = o_p(d^{-\frac{\varsigma}{2}}), \text{ where } \varsigma \in [0, \alpha - \eta - \theta].$$

Since $H_j = (h_{1,j}, \dots, h_{d,j})^T = \sum_{k=2}^d z_{k,j} u_k, j = 1, \dots, n$, it follows that $h_{i,j} = \sum_{k=2}^d u_{i,k} z_{k,j} = \sum_{k=2}^d u_{i,k} \lambda_k^{\frac{1}{2}} w_{k,j} \sim N(0, \sigma_{i,j}^2)$, where $\sigma_{i,j}^2 \leq \lambda_2, i =$

$1, \dots, [d^\beta]$, $j = 1, \dots, n$. Thus, for fix τ

$$\begin{aligned} P \left[\sum_{i=1}^{[d^\beta]} \sum_{j=1}^n d^{\frac{\varsigma}{2}} \lambda_1^{-\frac{1}{2}} |u_{i,1} h_{i,j}| \geq \tau \right] &\leq P \left[\bigcup_{i=1}^{[d^\beta]} \left\{ \sum_{j=1}^n |u_{i,1} h_{i,j}| \geq d^{-\frac{\varsigma}{2}} \lambda_1^{\frac{1}{2}} \tau u_{i,1}^2 \right\} \right] \\ &\leq \sum_{i=1}^{[d^\beta]} \sum_{j=1}^n P \left[|h_{i,j}| \geq n^{-1} d^{-\frac{\varsigma}{2}} \lambda_1^{\frac{1}{2}} \tau |u_{i,1}| \right] \leq \sum_{i=1}^{[d^\beta]} \sum_{j=1}^n P \left[|h_{i,j} \sigma_{i,j}^{-1}| \geq c^* d^{\frac{\alpha-\eta-\theta-\varsigma}{2}} \right] \\ &= 2n[d^\beta] \int_{cd^{(\alpha-\eta-\theta-\varsigma)/2}}^{+\infty} \frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{x^2}{2} \right\} dx \longrightarrow 0, \text{ as } d \rightarrow \infty, \end{aligned}$$

where c is constant. Similar, we can show that

$$(7.7) \quad \sum_{i=1}^{[d^\beta]} \lambda_1^{-1} \left(\sum_{j=1}^n |h_{i,j}| \right)^2 = o_p(d^{\frac{\varsigma}{2}}),$$

and

$$(7.8) \quad \sum_{i=[d^\beta]+1}^d \sum_{j=1}^n \lambda_1^{-\frac{1}{2}} |h_{i,j}| 1_{\{\sum_{j=1}^n |h_{i,j}| > \lambda\}} = o_p(d^{-\frac{\varsigma}{2}}),$$

where $\varsigma \in [0, \alpha - \eta - \theta)$.

Finally, we want to show that

$$(7.9) \quad \sum_{i=1}^{[d^\beta]} u_{i,1}^2 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}} = o_p(d^{-\frac{\varsigma'}{2}}),$$

where ς' satisfies that $d^{\frac{\varsigma'+\eta-\alpha}{2}} \lambda = o(1)$. Since we can always find a subsequence of $\{\frac{\lambda_1}{\sum_{i=2}^d \lambda_i}\}$ and make it convergent to a nonnegative constant, for simplicity, we just assume that $\lim_{d \rightarrow \infty} \frac{\lambda_1}{\sum_{i=2}^d \lambda_i} = C$. If $C = 0$, then the spike index $\alpha < 1$, and Jung and Marron (2009) [?] showed that

$$c_d^{-1} S_d \xrightarrow{p} I_n, \text{ as } d \rightarrow \infty,$$

where $c_d = n^{-1} \sum_{i=1}^d \lambda_i$. Since the eigenvector \tilde{v}_1^T of $c_d^{-1} S_d$ can be chosen so that they are continuous according to Acker (1974) [?], it follows that $\tilde{v}_1^T \Rightarrow v_1$, as $d \rightarrow \infty$, where \Rightarrow denotes the convergence in distribution and v_1 is the first eigenvector of n -dimensional identity matrix. If $C = 0$, then

the spike index $\alpha = 1$ and Jung, Sen and Marron (2010) [?] showed that $\tilde{v}_1^T \Rightarrow \frac{\tilde{W}_1}{\|\tilde{W}_1\|}$, as $d \rightarrow \infty$. Therefore, we have

$$(7.10) \quad |\tilde{v}_1^T \tilde{W}_1| \Rightarrow |v_1^T \tilde{W}_1| \text{ or } \|\tilde{W}_1\|, \text{ as } d \rightarrow \infty.$$

Since $d^{\frac{\zeta'+\eta-\alpha}{2}}\lambda = o(1)$, $d^{\frac{\zeta'+\eta-\alpha}{2}}\sum_{j=1}^n \max_{1 \leq i \leq [d^\beta]} |h_{i,j}| = o_p(1)$, and

$$\begin{aligned} & \sum_{i=1}^{[d^\beta]} d^{\frac{\zeta'}{2}} u_{i,1}^2 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}} \leq \sum_{i=1}^{[d^\beta]} d^{\frac{\zeta'}{2}} u_{i,1}^2 1_{\{|u_{i,1} \tilde{v}_1^T \tilde{Z}_1| \leq |\tilde{v}_1^T \tilde{H}_i| + \lambda\}} \\ & \leq \sum_{i=1}^{[d^\beta]} d^{\frac{\zeta'}{2}} u_{i,1}^2 1_{\left\{|\tilde{v}_1^T \tilde{W}_1| \leq \lambda_1^{-\frac{1}{2}} \max_{1 \leq i \leq [d^\beta]} |u_{i,1}|^{-1} \left(\sum_{j=1}^n \max_{1 \leq i \leq [d^\beta]} |h_{i,j}| + \lambda\right)\right\}} \\ & \leq \frac{cd^{\frac{\zeta'+\eta-\alpha}{2}} \sum_{j=1}^n \max_{1 \leq i \leq [d^\beta]} |h_{i,j}| + cd^{\frac{\zeta'+\eta-\alpha}{2}} \lambda}{|\tilde{v}_1^T \tilde{W}_1|}, \end{aligned}$$

where c is a constant, it follows that (7.9) is established.

Then, from (7.4), (7.5), (7.6), and (7.9), we obtain the following result about the numerator

$$(7.11) \quad \lambda_1^{-\frac{1}{2}} \left| \sum_{i=1}^{[d^\beta]} \check{u}_{i,1} u_{i,1} \right| = |\tilde{v}_1^T \tilde{W}_1| + o_p(d^{-\frac{\min\{\zeta, \zeta'\}}{2}}).$$

Similarly for the denominator, we have

$$\begin{aligned} (7.12) \quad & \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^d \check{u}_{i,1}^2} \leq \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^{[d^\beta]} \check{u}_{i,1}^2} + \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=[d^\beta]+1}^d \check{u}_{i,1}^2} \\ & \leq \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^{[d^\beta]} \left(u_{i,1} \tilde{v}_1^T \tilde{Z}_1\right)^2} + \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^{[d^\beta]} (\tilde{v}_1^T \tilde{H}_i)^2} + \lambda_1^{-\frac{1}{2}} \sum_{i=[d^\beta]+1}^d |\check{u}_{i,1}| \\ & = |\tilde{v}_1^T \tilde{W}_1| + \sqrt{\sum_{i=1}^{[d^\beta]} \lambda_1^{-1} \left(\sum_{j=1}^n |h_{i,j}|\right)^2} + \\ & + \sum_{i=[d^\beta]+1}^d \sum_{j=1}^n \lambda_1^{-\frac{1}{2}} |h_{i,j}| 1_{\{\sum_{j=1}^n |h_{i,j}| > \lambda\}}, \end{aligned}$$

and

$$\begin{aligned}
 (7.13) \quad \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^d \check{u}_{i,1}^2} &\geq \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^{\lfloor d^\beta \rfloor} \check{u}_{i,1}^2} \\
 &\geq |\tilde{v}_1^T \tilde{W}_1| - |\tilde{v}_1^T \tilde{W}_1| \sqrt{\sum_{i=1}^{\lfloor d^\beta \rfloor} u_{i,1}^2 1_{\{|\tilde{v}_1^T \tilde{X}_i| \leq \lambda\}}} - \sqrt{\sum_{i=1}^{\lfloor d^\beta \rfloor} \lambda_1^{-1} \left(\sum_{j=1}^n |h_{i,j}| \right)^2}.
 \end{aligned}$$

Combining (7.12), (7.13), (7.7), (7.8) and (7.9), we have

$$(7.14) \quad \lambda_1^{-\frac{1}{2}} \sqrt{\sum_{i=1}^d \check{u}_{i,1}^2} = |\tilde{v}_1^T \tilde{W}_1| + o_p(d^{-\frac{\min\{\varsigma, \varsigma'\}}{2}}).$$

Furthermore, (7.3), (7.10), (7.11), and (7.14) suggest that

$$| \langle \hat{u}_1^{\text{ST}}, u_1 \rangle | = \frac{|\tilde{v}_1^T \tilde{W}_1| + o_p(d^{-\frac{\min\{\varsigma, \varsigma'\}}{2}})}{|\tilde{v}_1^T \tilde{W}_1| + o_p(d^{-\frac{\min\{\varsigma, \varsigma'\}}{2}})} = 1 + o_p(d^{-\frac{\min\{\varsigma, \varsigma'\}}{2}}),$$

which means that \hat{u}_1^{ST} is consistent with u_1 with convergence rate $d^{-\frac{\min\{\varsigma, \varsigma'\}}{2}}$. This concludes the proof for Theorem 2.2.

In addition, note that $d^{\frac{\varsigma' + \eta - \alpha}{2}} \lambda = o(1)$. If $\lambda = o(d^{\frac{\alpha - \eta - \varsigma}{2}})$, then we can take $\varsigma' = \varsigma$. Then \hat{u}_1^{ST} is consistent with u_1 with convergence rate $d^{\frac{\varsigma}{2}}$. This finishes the proof of Theorem 2.3.

7.2. Proofs of Theorem 3.1, 3.2, 3.3, 3.4 and 3.5. The proofs of Theorems 3.1, 3.2, 3.4 and 3.5 are modifications of the proofs of Theorems 2.2 and 2.3. These are provided in the supplementary material, available online at [?]. The proof of Theorem 3.3 is also given in the supplement.

7.3. Proof of Theorem 4.1. Since $X_j^* = (I_{\lfloor d^\beta \rfloor}, (0)_{\lfloor d^\beta \rfloor \times (d - \lfloor d^\beta \rfloor)}) X_j$, where $I_{\lfloor d^\beta \rfloor}$ denotes the $\lfloor d^\beta \rfloor$ -dimensional identity matrix and $(0)_{\lfloor d^\beta \rfloor \times (d - \lfloor d^\beta \rfloor)}$ is the $\lfloor d^\beta \rfloor$ -by- $(d - \lfloor d^\beta \rfloor)$ zero matrix, $j = 1, \dots, n$, it follows that

$$\Sigma_{\lfloor d^\beta \rfloor}^* = (I_{\lfloor d^\beta \rfloor}, (0)_{\lfloor d^\beta \rfloor \times (d - \lfloor d^\beta \rfloor)}) \Sigma_d (I_{\lfloor d^\beta \rfloor}, (0)_{\lfloor d^\beta \rfloor \times (d - \lfloor d^\beta \rfloor)})^T,$$

which yields

$$\lambda_1^* = \lambda_1, \lambda_2 \geq \lambda_i^* \geq \lambda_d, j = 2, \dots, \lfloor d^\beta \rfloor.$$

Therefore,

$$(7.15) \quad \frac{\sum_{i=2}^{\lfloor d^\beta \rfloor} \lambda_i^{*2}}{(\sum_{i=2}^{\lfloor d^\beta \rfloor} \lambda_i^*)^2} \leq \frac{\lfloor d^\beta \rfloor \lambda_2^2}{(\lfloor d^\beta \rfloor)^2 \lambda_d^2} = \frac{O(d^\beta)O(d^{2\theta})}{O(d^{2\beta})} = o(1),$$

and

$$(7.16) \quad \frac{\lambda_1^*}{\sum_{i=2}^{\lfloor d^\beta \rfloor} \lambda_i^*} \leq \frac{\lambda_1}{\lfloor d^\beta \rfloor \lambda_d} = \frac{O(d^\alpha)}{O(d^\beta)} = o(1).$$

If we rescale λ_i^* , $i = 1, \dots, \lfloor d^\beta \rfloor$, (7.15) satisfies the ε_2 assumption of Jung and Marron (2009) [?] and (7.16) satisfies the assumption $\lambda_1 = O(d^\alpha)$ and $\sum_{i=2}^d \lambda_i = O(d)$, where $\alpha < 1$. For this case, Jung and Marron (2009) [?] have shown that \hat{u}_1^* is strongly inconsistent with u_1^* . This means that the oracle estimator \hat{u}_1^{OR} is strongly inconsistent with u_1 .

Acknowledgements. See [Supplement A](#) and [Supplement B](#) for the supplementary materials.

SUPPLEMENTARY MATERIAL

Supplement B: Additional Simulation Results

(<http://www.unc.edu/~dshen/RSPCA/B.html>). Additional simulation results are provided for the twenty spike index and sparsity index pairs, indicated in Figure 1.

DEPARTMENT OF STATISTICS AND OPERATIONS RESEARCH
UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL
CHAPEL HILL, NC 27599
E-MAIL: dshen@email.unc.edu
haipeng@email.unc.edu
marron@email.unc.edu